



# The Risc-V Journey Thru Containers to the Cloud

**Carlos Eduardo de Paula**


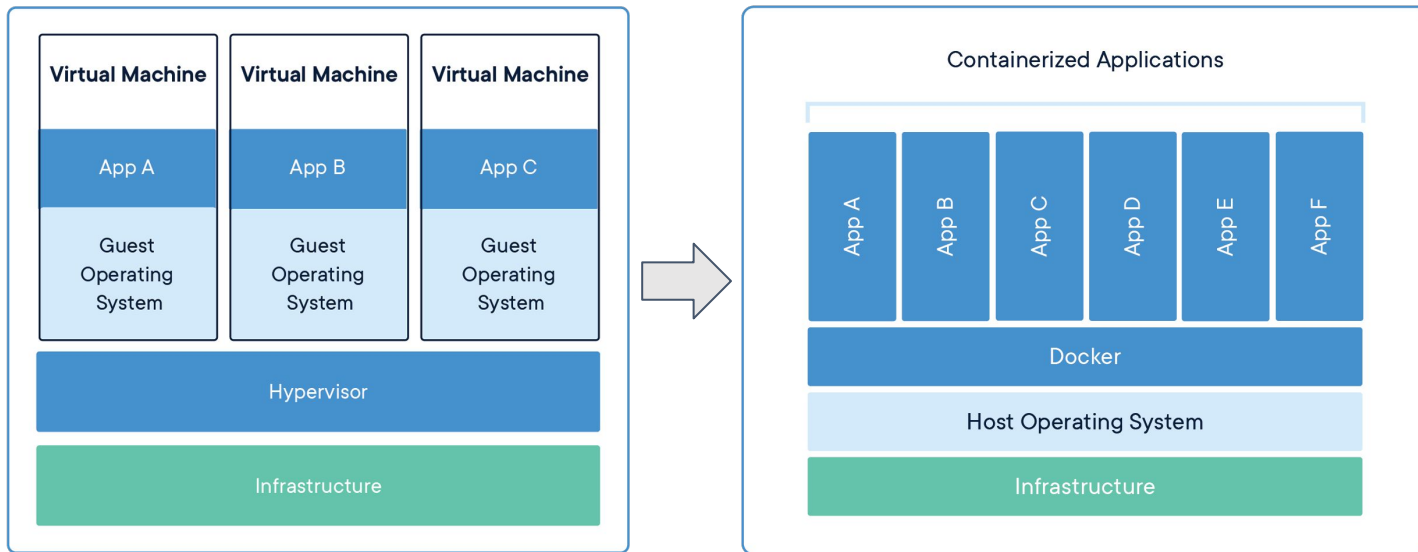
**Cloud Architect**

**Red Hat**

**@carlosedp**

#RISCVSUMMIT | [tmt.knect365.com/risc-v-summit/](https://tmt.knect365.com/risc-v-summit/)

# What are containers?



- **A standardized unit of software**
- **Packs all application requirements into a distribution image**
- **Runs everywhere. Developer laptop, QA server, production cloud with no changes**



# Can we have containers and cloud-native applications on Risc-V?

# Yes, we can!

docker hub

Search for great content (e.g., mysql)

Explore



carlosedp/echo-riscv ☆

By carlosedp • Updated

Container

Overview

Tags

Filter Tags

latest

Last updated 3 months ago by carlosedp

DIGEST

b23012a5adac

8b03ebeb8af8

5425959f37b5



```
~/echo
> uname -a
Linux qemuriscv 5.3.0-rc4-g6c2a8bd-dirty #6 SMP Fri Nov 22 18:08:43 -02 2019 riscv64 GNU/Linux

~/echo
> DOCKER_BUILDKIT=1 docker build -t carlosedp/echo-riscv .
[+] Building 0.3s (5/5) FINISHED
=> [internal] load .dockerignore
=> => transferring context: 2B
=> [internal] load build definition from Dockerfile
=> => transferring dockerfile: 37B
=> [internal] load build context
=> => transferring context: 34B
=> CACHED [1/1] COPY echo-riscv /echo-riscv
=> exporting to image
=> => exporting layers
=> => writing image sha256:231b741f511088f6f3901b8877ba8130f8e14a0780fa5ae261dd1f1c4d6b81f2
=> => naming to docker.io/carlosedp/echo-riscv

~/echo
> docker run --name hello -d -p 8085:8080 carlosedp/echo-riscv
bf363e2f2b19a7171c88664e685289882a167d00c3118bfe77138ef11aae9899

~/echo
> curl localhost:8085
Hello, I'm running Echo inside a container on linux/riscv64
```

#RISCVSUMMIT | [tmt.knect365.com/risc-v-summit/](https://tmt.knect365.com/risc-v-summit/)

# And Container Orchestration?

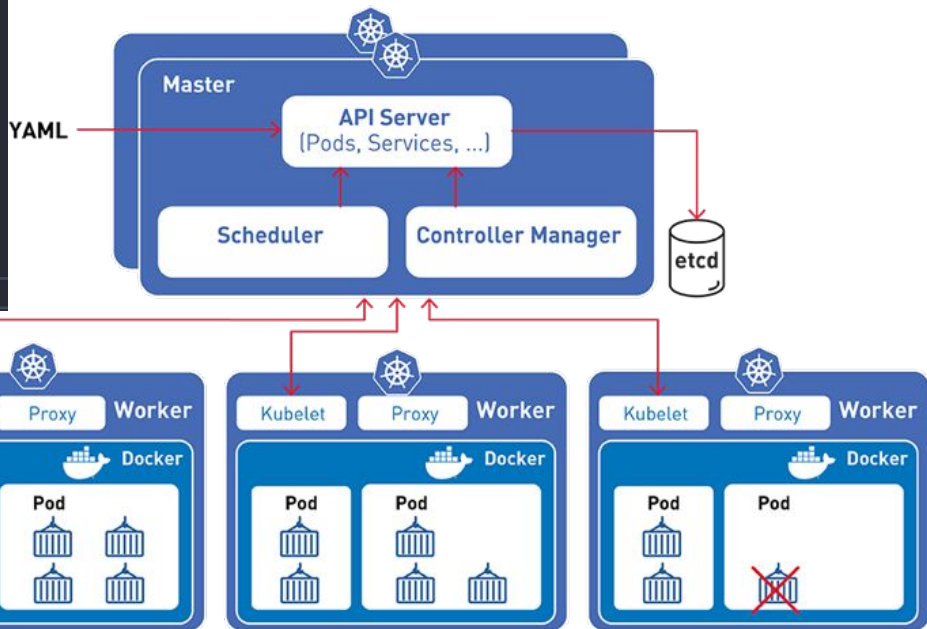


+

kubernetes



```
apiVersion: v1
kind: Pod
metadata:
  name: static-web
  labels:
    role: myapp
spec:
  containers:
  - name: web
    image: nginx
    ports:
    - name: web
      containerPort: 80
```



# Orchestration Works on Risc-V too!

```
Unleashed: ~
kubernetes-admin@kubernetes: ~ < carlosedp@unleashed
> echo "I'm OpenFaaS running on Kubernetes in Risc-V" | faas-cli invoke figlet
I'm OpenFaaS
running on
Kubernetes in
Risc-V

Last login: Sun Nov 24 18:11:05 2019 from 192.168.15.139
carlosedp@unleashed
-----
OS: Debian GNU/Linux bullseye/sid riscv64
Host: 5400e9c90806c
Kernel: 5.3.0-rc4-g6c2a8bd-dirty
Uptime: 3 days, 29 mins
Shell: zsh 5.7.1
Terminal: /dev/pts/3
CPU: (4) @ 1.400GHz
Memory: 2779MiB / 7973MiB

kubernetes-admin@kubernetes: ~ < carlosedp@unleashed
> uname -a
Linux unleashed 5.3.0-rc4-g6c2a8bd-dirty #5 SMP Fri Aug 16 18:28:38 -03 2019 riscv64 GNU/Linux

kubernetes-admin@kubernetes: ~ < carlosedp@unleashed

Unleashed: ~
> k get nodes -o wide
NAME          STATUS  ROLES  AGE   VERSION          INTERNAL-IP  EXTERNAL-IP  OS-IMAGE          KERNEL-VERSION  CONTAINER-RUNTIME
unleashed    Ready   master  65d   v0.0.0-master+... 192.168.15.16 <none>        Debian GNU/Linux bullseye/sid 5.3.0-rc4-g6c2a8bd-dirty docker://Unknown

Unleashed: ~
> kpod
NAMESPACE  NAME                                     READY  STATUS  RESTARTS  AGE   IP             NODE          NOMINATED NODE  READINESS GATES
default    echo-85c44d68bb-8bhf1                   1/1    Running  0          2d23h 10.244.0.175   unleashed      <none>           <none>
kube-system coredns-98fcf896-dnvtm                  1/1    Running  15         65d    10.244.0.169   unleashed      <none>           <none>
kube-system coredns-98fcf896-sn858                   1/1    Running  15         65d    10.244.0.170   unleashed      <none>           <none>
kube-system etcd-unleashed                          1/1    Running  4          65d    192.168.15.16  unleashed      <none>           <none>
kube-system kube-apiserver-unleashed                1/1    Running  5d21h     192.168.15.16  unleashed      <none>           <none>
kube-system kube-controller-manager-unleashed    1/1    Running  16         3d     192.168.15.16  unleashed      <none>           <none>
kube-system kube-flannel-ds-riscv64-g54xw       1/1    Running  0          8d     192.168.15.16  unleashed      <none>           <none>
kube-system kube-proxy-795ht                   1/1    Running  3          65d    192.168.15.16  unleashed      <none>           <none>
kube-system kube-scheduler-unleashed            1/1    Running  22         60d    192.168.15.16  unleashed      <none>           <none>
kube-system traefik-ingress-controller-5f844bb54d-xcw2d 1/1    Running  0          5d2h   10.244.0.163   unleashed      <none>           <none>
openfaas-fn figlet-54cd4964b5-b42jp                   1/1    Running  0          64d    10.244.0.171   unleashed      <none>           <none>
openfaas alertmanager-c467d66c4-htlbc        1/1    Running  0          60d    10.244.0.168   unleashed      <none>           <none>
openfaas basic-auth-plugin-ccc64f5f8-lclc8    1/1    Running  0          3d     10.244.0.173   unleashed      <none>           <none>
openfaas faas-idler-6fd4b449b5-vwkmf          1/1    Running  0          2d23h 10.244.0.174   unleashed      <none>           <none>
openfaas gateway-758569f6f5-xhswf           2/2    Running  1          3d     10.244.0.172   unleashed      <none>           <none>
openfaas nats-84794777cf-tdt2m                1/1    Running  0          64d    10.244.0.165   unleashed      <none>           <none>
openfaas prometheus-6df94c7df7-67fl2         1/1    Running  0          64d    10.244.0.164   unleashed      <none>           <none>
openfaas queue-worker-79859bd-cfkkf          1/1    Running  0          64d    10.244.0.167   unleashed      <none>           <none>
```



kubernetes



Prometheus



OPENFAAS

# Container Image list is growing!

## Base Container Images

- Debian Sid -> `carlosedp/debian:sid`
- Debian Sid Slim -> `carlosedp/debian:sid-slim`
- Busybox (1.31.0) -> `carlosedp/busybox:1.31`
- Go 1.13 (Multiarch) -> `carlosedp/golang:1.13`



## Prometheus

- Prometheus - `carlosedp/prometheus:v2.11.1-riscv64`
- AlertManager - `carlosedp/alertmanager:v0.18.0-riscv64`

## Traefik

- traefik v2 - `carlosedp/traefik:v2.0-riscv64`
- Whoami (Traefik demo) - `carlosedp/whoami:riscv64`



## Kubernetes


- kube-apiserver - `carlosedp/kube-apiserver:1.16.0`
- kube-scheduler - `carlosedp/kube-scheduler:1.16.0`
- kube-controller-manager - `carlosedp/kube-controller-manager:1.16.0`
- kube-proxy - `carlosedp/kube-proxy:1.16.0`
- pause - `carlosedp/pause:3.1`
- flannel - `carlosedp/flannel:v0.11.0`
- etcd (v3.5.0) - `carlosedp/etcd:3.3.10`
- CoreDNS (v1.6.3) - `carlosedp/coredns:v1.6.2`

Kubernetes images are multi-arch with manifests to `arm`, `arm64`, `amd64`, `riscv64` and `ppc64le`.

## Misc Images

- Echo demo - `carlosedp/echo-riscv`
- CoreDNS v1.3.0 - `carlosedp/coredns:v1.3.0-riscv64`

# How we got there? Golang first!



- First working version around may/19
- Passes all tests
- Expected upstreaming in 1.14

all: port to RISC-V #27532

**Open** tklausner opened this issue on Sep 6, 2018 · 165 comments

Tree kept in:

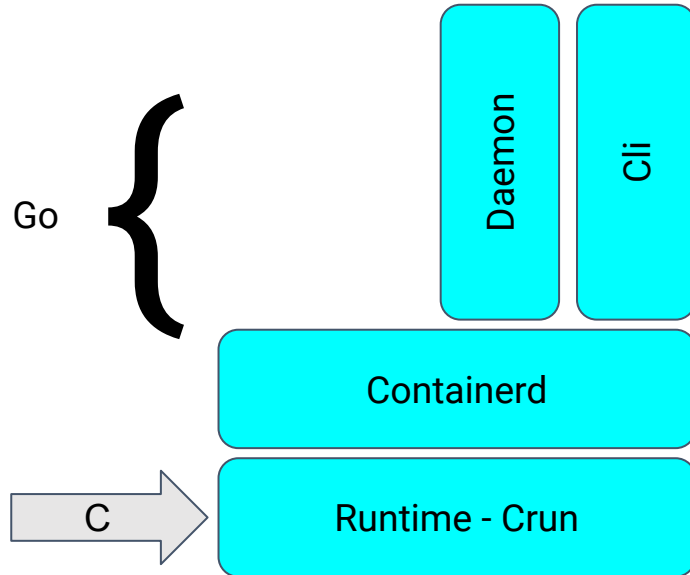
<https://github.com/4a6f656c/riscv-go>

Subject	Status	Owner	Assignee	Repo	Branch	Updated	Size	CR	RT	TI
☆ cmd/internal/obj/riscv: split immediates larger than 12-bits	Merged	Joel Sing	–	go	master	3:10 AM	M	✓	✓	✓
☆ cmd/link/cmd/internal/obj/riscv: implement archreloc for riscv64	Merged	Joel Sing	–	go	master	3:08 AM	M	✓	✓	✓
☆ cmd/compile: implement compiler for riscv64	Merge Conflict	Joel Sing	–	go	master	Nov 15	XL			
☆ cmd/internal/obj/riscv: implement prologue and epilogue	Merge Conflict	Joel Sing	–	go	master	Nov 15	L			
☆ cmd/internal/obj/riscv: handle AGETCALLERPC	–	Joel Sing	–	go	master	Nov 15	S	✓		
☆ cmd/internal/obj/riscv: handle FEQ/FNEG/SEQZ/SNEZ	Merge Conflict	Joel Sing	–	go	master	Nov 15	M			
☆ cmd/internal/obj/riscv: handle call, jmp and branch	–	Joel Sing	–	go	master	Nov 15	XL			
☆ runtime: add support for linux/riscv64	–	Joel Sing	–	go	master	Nov 13	XL			
☆ cmd/internal/objabi,cmd/link: initial linker support for riscv64	Merged	Joel Sing	–	go	master	Nov 13	M	✓	✓	✗
☆ internal/cpu/internal/bytealg: add support for riscv64	Merged	Joel Sing	–	go	master	Nov 12	M	✓	✓	✓
☆ [dev.link] Merge branch 'dev.link' of https://go.googlesource.com/go into dev.link	Abandoned	Jeremy Faller	–	go	master	Nov 07	XL			
☆ reflect: add support for riscv64	–	Joel Sing	–	go	master	Nov 06	S			
☆ math: add support for riscv64	–	Joel Sing	–	go	master	Nov 03	M			
☆ internal/syscall/unix: add support for linux/riscv64	Merged	Tobias Klausner	–	go	master	Nov 01	S	✓	✓	✓
☆ syscall: add support for linux/riscv64	Merged	Tobias Klausner	–	go	master	Nov 01	XL	✓	✓	✓
☆ cmd/internal/obj/riscv: implement MOV pseudo-instructions	Merged	Joel Sing	–	go	master	Oct 29	L	✓	✓	✓
☆ [dev.link] all: merge branch 'master' into dev.link	Merged	Cherry Zhang	–	go	dev.link	Oct 18	XL	✓	✓	✓

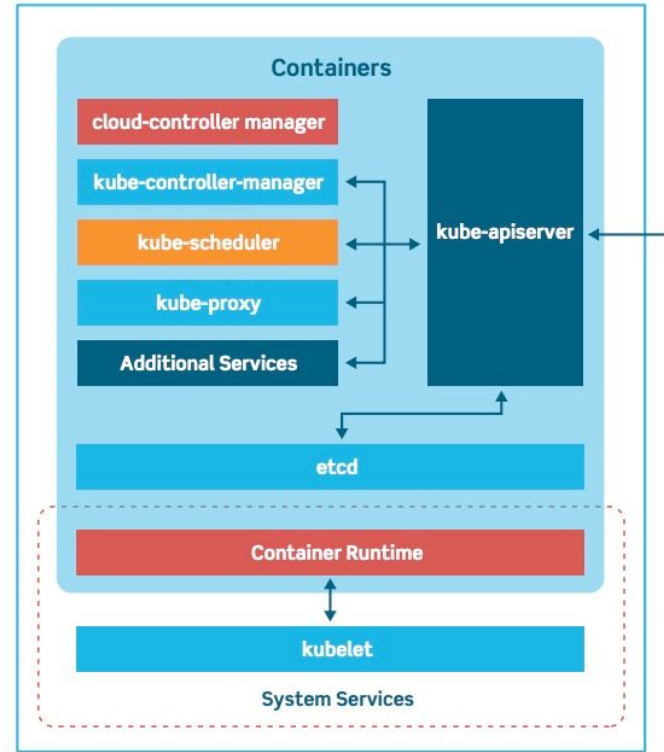


# Container Runtimes / Orchestrators

Docker / Podman



Kubernetes



# Some Pull Requests later

Project	Description	Pull Request	Upstream
go/net	ipv4, ipv6, internal/socket: add riscv64 support	<a href="https://github.com/golang/net/pull/43">https://github.com/golang/net/pull/43</a>	✓
go/sys	unix: add riscv64 tag to endian_little.go	<a href="https://github.com/golang/sys/pull/38">https://github.com/golang/sys/pull/38</a>	✓
mattn/go-isatty	Update x/sys to support Risc-V	<a href="https://github.com/mattn/go-isatty/pull/39">https://github.com/mattn/go-isatty/pull/39</a>	✓
creack/pty	Add riscv64 support	<a href="https://github.com/creack/pty/pull/81">https://github.com/creack/pty/pull/81</a>	✓
docker/cli	Add riscv64 to manifest annotation and bash completion	<a href="https://github.com/docker/cli/pull/2084">https://github.com/docker/cli/pull/2084</a>	✓
opencontainers/runc	Bump x/sys and update syscall for initial Risc-V support	<a href="https://github.com/opencontainers/runc/pull/2123">https://github.com/opencontainers/runc/pull/2123</a>	✓
moby/moby	bump x/sys to fix riscv64 epoll	<a href="https://github.com/moby/moby/pull/39726">https://github.com/moby/moby/pull/39726</a>	✓
moby/moby	Update modules to support riscv64	<a href="https://github.com/moby/moby/pull/39423">https://github.com/moby/moby/pull/39423</a>	✓
moby/moby	allow dockerd builds without cgo - Tonis	<a href="https://github.com/moby/moby/pull/39327">https://github.com/moby/moby/pull/39327</a>	✓
containerd/containerd	bump x/sys to fix riscv64 epoll	<a href="https://github.com/containerd/containerd/pull/3526">https://github.com/containerd/containerd/pull/3526</a>	✓
containerd/containerd	Update x/sys, x/net and bbolt modules to support Risc-V architecture	<a href="https://github.com/containerd/containerd/pull/3328">https://github.com/containerd/containerd/pull/3328</a>	✓
docker/cli	bump x/sys to fix riscv64 epoll	<a href="https://github.com/docker/cli/pull/2042">https://github.com/docker/cli/pull/2042</a>	✓
docker/libnetwork	bridge: add riscv64 build tags - Tonis	<a href="https://github.com/docker/libnetwork/pull/2389">https://github.com/docker/libnetwork/pull/2389</a>	✓
LK4D4/vndr	add riscv64 support	<a href="https://github.com/LK4D4/vndr/pull/80">https://github.com/LK4D4/vndr/pull/80</a>	✓
vishvananda/netns	add riscv64 architecture	<a href="https://github.com/vishvananda/netns/pull/34">https://github.com/vishvananda/netns/pull/34</a>	✓
containers/libpod	build: allow to build without cgo on RISC-V	<a href="https://github.com/containers/libpod/pull/3437">https://github.com/containers/libpod/pull/3437</a>	✓

## And many more...

Project	Description	Pull Request	Upstream
kubernetes/kubernetes	Bump x/sys and opencontainers/runc to support Risc-V architecture	<a href="https://github.com/kubernetes/kubernetes/pull/82349">https://github.com/kubernetes/kubernetes/pull/82349</a>	✗
kubernetes/kubernetes	Bump gonvml module and remove CGO dependency	<a href="https://github.com/kubernetes/kubernetes/pull/82342">https://github.com/kubernetes/kubernetes/pull/82342</a>	✓
etcd-io/etcd	vendor: update x/sys and x/net modules to support Risc-V	<a href="https://github.com/etcd-io/etcd/pull/10000">https://github.com/etcd-io/etcd/pull/10000</a>	✓
prometheus/prometheus	Riscv support	<a href="https://github.com/prometheus/prometheus/pull/10000">https://github.com/prometheus/prometheus/pull/10000</a>	✓
prometheus/promu	Update x/sys and x/net modules to support Risc-V	<a href="https://github.com/prometheus/promu/pull/10000">https://github.com/prometheus/promu/pull/10000</a>	✓
prometheus/alertmanager	Update x/sys and x/net modules to support Risc-V	<a href="https://github.com/prometheus/alertmanager/pull/10000">https://github.com/prometheus/alertmanager/pull/10000</a>	✓
containous/traefik	Bump x/sys to support Risc-V architecture	<a href="https://github.com/containous/traefik/pull/10000">https://github.com/containous/traefik/pull/10000</a>	✓
nats-io/nats-streaming-server	Update x/sys and etcd/bbolt modules to support Risc-V	<a href="https://github.com/nats-io/nats-streaming-server/pull/10000">https://github.com/nats-io/nats-streaming-server/pull/10000</a>	✓
openfaas/faas-swarm	Update x/sys and x/net modules to support Risc-V	<a href="https://github.com/openfaas/faas-swarm/pull/10000">https://github.com/openfaas/faas-swarm/pull/10000</a>	✓
alexellis/inlets	Update x/sys and x/net modules to support Risc-V	<a href="https://github.com/alexellis/inlets/pull/10000">https://github.com/alexellis/inlets/pull/10000</a>	✓
openfaas/faas-cli	Update x/sys module to support Risc-V	<a href="https://github.com/openfaas/faas-cli/pull/10000">https://github.com/openfaas/faas-cli/pull/10000</a>	✓
debuerreotype/debuerreotype	Add support to debian ports by passing --port	<a href="https://github.com/debuerreotype/debuerreotype/pull/10000">https://github.com/debuerreotype/debuerreotype/pull/10000</a>	✓
labstack/gommon	Update x/sys module to support Risc-V	<a href="https://github.com/labstack/gommon/pull/10000">https://github.com/labstack/gommon/pull/10000</a>	✓
labstack/echo	Update x/sys and x/net modules to support Risc-V	<a href="https://github.com/labstack/echo/pull/10000">https://github.com/labstack/echo/pull/10000</a>	✓
github/hub	Update x/sys module to support Risc-V	<a href="https://github.com/github/hub/pull/10000">https://github.com/github/hub/pull/10000</a>	✓
google/go-jsonnet	Update x/sys to a newer version that supports Risc-V architecture	<a href="https://github.com/google/go-jsonnet/pull/284">https://github.com/google/go-jsonnet/pull/284</a>	✓

**20+ Projects**  
**30+ PRs**  
**6 CNCF Projects**  
**95%+ Upstream**

# Still Pending

Project	Description	Status
Golang	Upstreaming language support for Risc-V	Planned for Go 1.14. Build box already been set.
SECCOMP	Add Risc-V Seccomp support to Linux Kernel	Planned for Kernel 5.5
Libseccomp	Upstream libseccomp patch	Pending Kernel SECCOMP support
Kubernetes bump	Update dependencies for upstream	Planned for Kubernetes 1.17
Golang CGO	Go extension for importing C code	Pending implementation
Runc	Container Runtime in C	Depends on CGO. CRUN is an alternative now.
Golang SQLite	File database libraries	Depends on CGO
Official Images	Build and publish official images	Depends on distro support (Debian, Alpine, Fedora)
Most projects	Build and distribute binaries and/or images	Changes on CI/CD and official images



# Thanks!

Contact me:

<https://twitter.com/carlosedp>

Read me:

<https://medium.com/@carlosedp>

Follow up progress on:

<https://github.com/carlosedp/riscv-bringup>

<https://hub.docker.com/u/carlosedp>